

What is claimed is:

1. A computer system comprising:
a network;
one or more processing nodes connected via the network, wherein each processing node includes:
a plurality of processors, wherein each processor includes a scalar processing unit, a vector processing unit and means for operating the scalar processing unit independently of the vector processing unit; and
a shared memory connected to each of the processors, wherein the shared memory includes a cache;
wherein processors on one node can load data directly from and store data directly to shared memory on another processing node via the network.
2. The computer system of claim 1, wherein the shared memory further includes a Remote Address Translation Table (RTT), wherein the RTT translates memory addresses received from a first processing node into physical addresses within the shared memory of a second processing node.
3. The computer system of claim 1, wherein the shared memory further includes a plurality of cache coherence directories, wherein each processing node is coupled to one of the cache coherence directories.
4. The computer system of claim 1, wherein each processor includes two vector pipelines.
5. The computer system of claim 1, wherein the processing nodes include at least one input/out (I/O) channel controller, wherein each I/O channel controller is coupled to the shared memory of the processing node.

6. The computer system of claim 1, wherein each scalar processing unit contains a scalar cache memory, wherein scalar cache memory contains a subset of cache lines stored in the shared memory cache.
7. The computer system according to claim 1, wherein the network includes a router connecting one or more of the processing nodes.
8. A computer system comprising:
 - a network;
 - one or more processing nodes connected via the network, wherein each processing node includes:
 - four processors configured as a Multi-Streaming Processor, wherein each processor includes a scalar processing unit, a vector processing unit and means for operating the scalar processing unit independently of the vector processing unit; and
 - a shared memory connected to each of the processors, wherein the shared memory includes four cache memories, wherein each cache memory is connected to each processor;
 - wherein processors on one node can load data directly from and store data directly to shared memory on another processing node via the network.
9. The computer system of claim 8, wherein the shared memory further includes a Remote Address Translation Table (RTT), wherein the RTT translates memory addresses received from a first processing node into physical addresses within the shared memory of a second processing node.
10. The computer system of claim 9, wherein the shared memory further includes a plurality of cache coherence directories, wherein each processing node is coupled to one of the cache coherence directories.

11. The computer system of claim 8, wherein the shared memory further includes a plurality of cache coherence directories, wherein each processing node is coupled to one of the cache coherence directories.